Exhibit D

The Markup

How We Built a Meta Pixel Inspector

By Surya Mattu, Angie Waller, Simon Fondrie-Teitler, and Micha Gorelick

April 28, 2022 08:00 ET

Viewable online at
https://themarkup.org/show-your-work/2022/04/28/how-we-built-a-meta-pixel-inspector

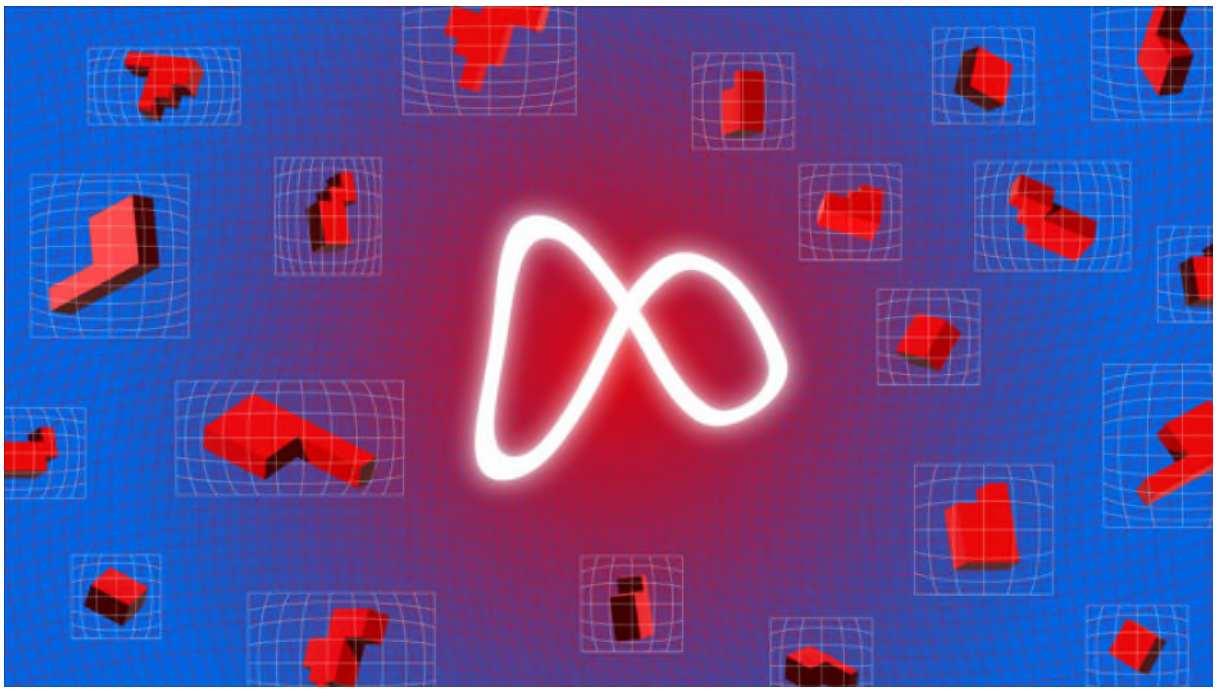

Introduction

Have you ever shopped for a product online only to see the item appear ad nauseam in your Facebook or Instagram feeds? This is often the result of the Meta Pixel (formerly called the Facebook Pixel), a sophisticated snippet of computer code used for tracking you around the web, embedded in the HTML of a website. This code works by sending Meta, Facebook's parent company, a detailed log of such user interactions as clicking on a product. The pixel's data collection is not limited to e-commerce: Even websites processing sensitive information (defined below), such as online gambling apps, can have the pixel installed. Few people are aware of how expansively Meta tracks their activities. Through Meta Pixel, for instance, a gambling app might notify Meta that you've registered with a particular email address—whether you have a Facebook account or not. The same tracking can occur as you submit an application for a student loan.

The Markup's Facebook Pixel Hunt, in collaboration with Mozilla Rally, is the **first large-scale crowdsourced study of the presence of the Meta Pixel and the data it collects in real-world scenarios**—when it is encountered while logging in to websites, submitting forms, buying products, and during everyday browsing activities. Using data collected from participants in our study, we monitor the pervasiveness of the Meta Pixel on the internet and the kind of information the tracker collects. As of April 25, we had 5,447 participants, but that number changes daily. A majority of participants (80 percent) have encountered the Meta Pixel on websites they visit. Sensitive user information might be sent to Meta from any number of sites without the person's awareness or consent. Our goal is to gain insight into that data collection. We describe our collaboration with Rally in the Our Study section. To learn about the technical details of how we assess whether Meta is collecting personally identifiable or sensitive information from a user, refer to the How We Analyze the Pixel Data section of this methodology.

The Meta Pixel gets its name from trackers that traditionally took the form of small, one-pixel-by-one-pixel images. These tiny graphics are embedded on websites and emails and typically collect info on who views the content. Since the Meta Pixel's first iteration over a decade ago, when it was called the Facebook Conversion Pixel, the pixel's functionality and tracking have grown quite expansive. Now the Meta Pixel is a mechanism that loads JavaScript code capable of collecting detailed and granular data for every interaction on a page. With all of this complexity, referring to it as only a "pixel" can be misleading.

See the data for our companion story here.

Meta can connect website visitors to their Facebook profile using third-party cookies. When a user logs into Facebook, Meta sets certain cookies on their browser. While this logged-in user browses other websites that contain the Meta Pixel, the tracker communicates with Meta's servers. When that happens, many web browsers—those that don't block third-party cookies—will also attach those previously set cookies. Like other third-party cookies, those set by the pixel allow Meta to build detailed dossiers about the site's users as those users traverse the web, so advertisers can target people with customized ads on Facebook and Instagram based on their online behavior. In case the cookies aren't enough to match a user browsing a website to a Facebook or Instagram profile, Meta also allows the website to send personal information a user enters in a form to match them to their Facebook or Instagram profile, even if they are not logged in to Facebook at the time. This feature is called Advanced Matching and is described in more detail in the Advanced Matching Parameters section.

The Meta Pixel collects data on website visitors regardless of whether or not they have Facebook or Instagram accounts. It is unclear what Meta does with that data for nonusers, but when asked by Congress about Meta, then Facebook, maintaining "shadow profiles" on nonusers of Facebook, Mark Zuckerberg responded that the company holds data on nonusers for "security purposes" such as preventing third-party collection of data.

Meta did not respond to The Markup's requests for comment on this methodology.

Background

Like other ad trackers, the Meta Pixel is installed and configured by website owners to measure things like user engagement on their sites and to target ads to their website visitors—all things the website owner can do through Meta's products. Blacklight, our real-time web privacy inspector, found that more than 30 percent of popular websites embed the Meta Pixel. In response to congressional questioning in 2018, Facebook said that more than two million of its pixels have been installed on websites.

Previous research has examined the widespread presence of Meta Pixels on the web, including in gambling apps, where the data collected could include frequency and types of games played or amount of cash deposited in an account, details that another advertiser or app could use to target people who have demonstrated compulsive gambling behaviors. Risk assessments have shown how website owners could use data collected through Meta Pixel to de-anonymize visitors to their

website. Other reports have shown that despite age restrictions, Meta Pixel tracking is not blocked for Facebook users under 16 years old.

In response to criticism following the Cambridge Analytica scandal, Facebook implemented new transparency tools so people can see what web and app activities outside of Facebook have sent behavioral data to their profile. After clicking through multiple menus on the privacy settings, a user can see Facebook's "off-Facebook activity," which shows which sites have shared information about that person with Meta. This only offers a snapshot, however, since Facebook does not disclose what information was collected or shared.

Prior Work

Other research initiatives use crowdsourcing by volunteers to capture evidence of Facebook's data surveillance for ad targeting.

- NYU Ad Observatory is a browser-extension-enabled project that archives and shares ads and metadata from Facebook's and Google's political ad libraries as well as targeted ads served to volunteers who've downloaded the extension and signed in to Facebook on their desktops.
- What Facebook Knows About You includes a browser extension that allows users to share with ProPublica all the "interests" Facebook has tagged them with—information that is typically only available to individual users.
- The Markup's Citizen Browser is a desktop application that collects and redacts data from the Facebook feeds of paid volunteers. Reporting from this project has explored Facebook's proxy interests, which allow targeting of users by race and partisan politics.
- The Markup's Blacklight tool can determine if the Meta Pixel is present on a given webpage. While it surveys websites, it does not examine the data that is sent through the pixel nor the prevalence of pixels on websites requiring login.

While these tools and datasets capture evidence of Meta Pixel's presence on websites and the kinds of pervasive ad targeting made possible by Meta's data collection, the Facebook Pixel Hunt is the first study to investigate the types and breadth of personal data that can be sent through Meta Pixel as people browse the web and enter information into websites.

Our Study

Rally Platform

In order to reliably and safely collect data while protecting user privacy, we partnered with Mozilla Rally. Rally is a data sharing research platform developed in 2021 by Mozilla in collaboration with researchers at Princeton University. Rally aims to counter the power imbalance of large companies controlling access to data and builds tools that empower users to gain insights into how their data is collected and used across the internet. The Rally software runs as a browser extension for Firefox that anyone can install and run to participate in public interest studies. When people join Rally studies, they are informed upfront about the nature of the study and what data will be collected. All Rally users consent to participation.

The Markup's Facebook Pixel Hunt study with Rally launched in January 2022. The goal of the study is to monitor Meta's pixel tracking mechanism and understand the kinds of information it collects on sites across the web. The study will continue until July 13, 2022. If you would like to learn more about the study or participate, visit the Facebook Pixel Hunt page on Rally. Additional information about the Rally platform and how it works is available on Mozilla's website. More details on how Rally protects user data can be found in its privacy policy.

Data Collection

Panel

> **Help Us Investigate Facebook Pixel Tracking**
>
> The full extent of how Facebook tracks people on the web is not understood—but you can help uncover what data the tech giant is gathering about you. Join our study.

To join the study, volunteers are asked to install the Rally browser add-on, which will detect the presence of the Meta Pixel as they browse the web. Mozilla then collects all the data that is shared with Meta through these pixel instances. The code for the browser add-on is open source. The data collected during the study is described in the metrics and pings configuration files in Mozilla's GitHub code repository.

Participants can opt in to a demographic survey as part of the study. As of April 25, 2022, 4,833 of 5,447 participants had completed at least part of this survey. Of those respondents who self-reported gender, a majority are male (79 percent). The majority of respondents who reported education level have a bachelor's degree or higher (63 percent), and of those who self-disclosed race, a majority are White (78 percent). Age ranges are broader, with the largest categories reported being 25 to 34 (24 percent) and 35 to 44 (20 percent). For a detailed breakdown, see the Panel Demographics section. For the limitation posed by the demographics, see the Limitations section.

Pixel Data We Capture

The browser add-on captures network requests that the pixel makes to Meta's server. These requests contain information about the user's interaction with the site. The data is sent either in the URL parameters of a GET request or the body of a POST request to the URL https://www.facebook.com/tr/. For more information about all the different types of data the pixel sends, please refer to Data Meta Collects in the appendix.

Meta Pixel "events" are the basic unit of analysis of this study. These are the packets of data the pixel code sends to Meta's servers from a website. Once an advertiser installs the pixel, by default it begins sending the "PageView," "Microdata" and "SubscribedButtonClick" events automatically.

The "PageView" event sends data that a person landed on the webpage. The "Microdata" event sends all of the meta information about the webpage pulled from the schema.org or OpenGraph tags and can include information like page title, URL, and page description. Meta does not provide documentation of what the "SubscribedButtonClick" event collects and only refers to it as "button click data" in the automatic configuration description. We have found that the "SubscribedButtonClick" event tracks form submission clicks, button clicks, and clicks on other elements of the page, including some links, not just clicks on subscribe buttons as the name suggests.

The advertiser can also add code to send additional events in reaction to the user's activities on the site. Meta broadly places these actions in two categories: standard events and custom events.

Standard events are actions that have been predefined by Meta and are recognized across all of its ad products. As of this writing, there are 18 standard events described in the company's documentation. These events describe common web

interactions such as adding an item to a cart, purchasing an item, submitting an application, donating, or subscribing to a newsletter. The PageView event that is sent by default on pixel initialization is also listed as a standard event.

Custom events fall outside those covered by standard events. Custom events are commonly used by the ecosystem of third-party social media analytics services and plug-ins that integrate with the Meta Pixel; these tools rely on custom events for some of their services. We have observed a range of custom events in our dataset, including how far down a page a user has scrolled, whether a user has reset their password, if a user viewed the video on the page, and if a user clicked on a certain part of the page such as a "more info" toggle.

As of April 25, we have seen a total of 2,635,130 pixel events, the majority of which (the first three noted below) are events sent by default:

- 39.1% (1,032,218) are PageView events.
- 22.8% (602,166) are Microdata events.
- 14.6% (384,286) are SubscribedButtonClick events.
- 9.0% (236,870) are standard events other than the default PageView.
- 14.5% (381,919) are custom events, including 5,520 unique custom event types.

How We Analyze the Pixel Data

We primarily focus on pixel parameters, described below, that may contain sensitive information. "Sensitive data" is any data that Meta or a website's owners could potentially use to identify a user or data that Meta states the owners "reasonably should know" is sensitive, such as information about children under the age of 13, health information, financial information, or similar categories of information.

Standard Parameters

When a website sends an event to Meta, it sends the event name with additional details. These details can include the URL where the event was triggered, the user's screen dimensions, the version number of the pixel, and technical details beyond the focus of our study.

What We Look For

We look in the Rally data for whether the website's pixel is sending sensitive data with the request. The Meta Pixel allows website owners to track a visitor's actions on a given page. Sometimes these actions happen on sites that require users to log in. This can lead to sensitive information such as the visitor's user name being sent with the request.

Custom Data Parameters

The pixel allows the addition of Custom Data Parameters for website owners to collect additional details about their visitors. These parameters can contain any data that pixel customers want to track about their site visitors. Typical examples include what content visitors viewed on the site or how far down the page they scrolled.

Meta itself also relies on these parameters to send additional data in events sent from the Meta Pixel by default (as opposed to set up by the website owner). For example, the Microdata event sends metadata about the site using these parameters. Third-party services such as customer data platforms and page analytics services have Facebook integrations that also create custom events.

What We Look For

We search for event data in the URL query parameters of GET requests and the body of the POST requests, with the key matching the format cd[EVENT_NAME] (for example "cd[buttonText]"), and determine whether they are sending personally identifiable information (PII) or sensitive information to Meta. Event Data Details in the appendix describes the data as seen in the network requests.

Advanced Matching Parameters

Advanced Matching Parameters allow Meta to connect collected event data to users, even if they do not have Facebook's browser cookies. If a website visitor is logged in to Facebook on their browser, many browsers will send that cookie data to Meta's servers along with other data the Meta Pixel sends. This allows Meta to connect the visitor to their Facebook profile. If the person is not logged in to Facebook, or they use a browser like Safari or Firefox that blocks Facebook's third-party cookies by default, Meta may still be able to connect the data to the person's account. Meta calls this feature "Advanced Matching." With Advanced Matching, personal information such as name, email, gender, address, or birth date that a site knows about a visitor or that is entered in forms on the website can be collected and used to connect other event data to the Facebook user. See the appendix to learn more

about the different types of Advanced Matching and what data Meta can collect through this technique.

What We Look For

We search for event data in the URL query parameters of GET requests and the body of the POST requests, with the key matching the format udff [PERSONAL_IDENTIFIER]. Meta's documentation provides a table with the personal identifiers that can be used for matching. For example, the key "udff[em]" contains an email address as its value. For a full list of available identifiers, please refer to the appendix.

Meta states in its documentation that tracking data used for Advanced Matching is hashed before it is sent using the SHA-256 algorithm to protect user privacy. However, researchers have shown that merely hashing user data does not truly protect it. Hashing does not prevent Meta from using this data to match people who visit websites to their Facebook profiles. For more information about exactly what personal information is sent, please refer to the Data Meta Collects section of the appendix.

It's worth noting that we can easily look for data a user has typed into a form being sent to Meta for the purposes of Advanced Matching because this behavior is standard and described in its documentation. We cannot test if a website is using custom code to grab data that users enter into their website and sending it to Meta via custom data parameters, since the structure of that data would vary site-by-site. There is no way for us to know what a user types into form fields on the site, so we're not able to automatically compare that data to data being sent to Meta.

Panel Demographics

When Rally participants install the Rally browser add-on, they are invited to complete an optional demographic survey. As of April 25, 4,833 participants (89 percent) had completed at least part of his survey. The following tables summarize the participants' self-disclosed demographic characteristics, with percentages measured by total respondents for each optional question.

What Happens to Data Collected

While our study is able to identify the types of data collected by the Meta Pixel, we do not fully know how the data is processed and used by the company. According to Facebook's Business Tools Terms of Use, the personal information that is sent as a

part of the Advanced Matching parameters is used for associating tracking data with a user's account. For more information on how we identify Advanced Matching parameters, please see the How We Analyze the Pixel Data section.

Through matching, Meta creates Custom Audiences, which save these users for the advertiser to later send targeted ads. This data also contributes to Lookalike Audiences, which allow advertisers to target people "similar" to their custom audience based on demographic and unspecified behavioral information. Meta's terms do not address how it processes other personal data, such as login details appearing in URLs, that might be inadvertently collected through Meta Pixel event data. Refraining from collecting and sharing data from children under the age of 13, health information, financial data, and other sensitive categories as determined by legal regulations is a responsibility Meta places on the pixel customer. Meta claims in its terms of service that it can suspend a Business Tools account if these rules are violated. It does not detail how it monitors and enforces these rules, nor if it has a protocol for deleting data that is shared in violation of these policies.

It is unclear how Meta processes custom event data that is configured by the advertiser, such as form fields that collect data not specified in the standard events. Meta only states it may use event data to support the objectives of a pixel customer's ad campaign, improve the effectiveness of ad delivery models, and determine the relevance of ads to people.

Limitations

We Can't Test Every Observed Event

When a website shows that a pixel event was triggered, we perform an additional inspection of the site to document the presence of the Meta Pixel code. In some cases, for instance when a user must have an account or be logged in to a site, we may not be able to access the URL where the pixel event recorded in our dataset occurred.

While we are able to document sensitive data in URL and form data that gets sent to Meta through pixel configurations, we do not know how or if the data gets processed by Meta.

Panel Demographics

Our study requires users to browse the internet on a desktop computer, use the Firefox browser, and voluntarily participate in the Rally study. As of this writing our panel skews heavily toward men and is not representative of U.S. demographics.

This limits our ability to make claims about the influence of the pixel on the overall U.S. population using solely panel data.

Ad Blockers and Tracking Protection

If a panelist is using a tracker blocker such as uBlock Origin or EFF's Privacy Badger, these tools will likely block the Meta Pixel's network requests, thus preventing the pixel from sending data to Meta and thus preventing us from observing their pixel events.

A similar limitation would occur if a panelist uses Mozilla Firefox's Facebook Container extension. This add-on silos Facebook activity and blocks network requests made by the Meta Pixel on sites other than Facebook itself, preventing Meta from associating information about a user's activity on websites outside Facebook to that user's Facebook identity.

Our study cannot determine if a participant has Facebook Container or other ad blocking resources installed. Taking this into consideration, our study might be undercounting the actual exposure of panelists to the Meta Pixel on the sites they visit.

Browser-Based Network Requests Are Not the Only Way Websites Share Data with Meta

Meta offers an API that allows advertisers to send pixel data about visitors to Meta directly from their website's servers rather than the user's browser. Pixel data that is collected and sent in this way won't be seen by the Rally browser extension. Therefore, this study may be undercounting the number of sites integrated with the Meta Pixel or miss some personal data being sent to Meta by the sites via the Conversions API.

Appendix

Advanced Matching Parameters

Advanced Matching is the term Meta uses for associating tracking data with a user's account. There are two types of Advanced Matching:

Manual Advanced Matching: The website developer manually decides which pieces of data to send via the Meta Pixel.

Automatic Advanced Matching: The Meta Pixel will automatically look for recognizable form fields and other sources on the website that contain information such as first name, surname, and email address. The Meta Pixel receives information entered in these fields along with the event, or action, that took place. The table below shows the format of each field that can be sent with this method.

Meta			
User Data	**Parameter**	**Format**	**Example**
Email	em	Unhashed lowercase or hashed SHA-256	jsmith@example.com or 6e3913852f512d76acff15d1e402c7502a5bbe6101745a7120a2a4833ebd2350
First Name	fn	Lowercase letters	john
Last Name	ln	Lowercase letters	smith
Phone	ph	Digits only including country code and area code	16505554444
External ID	external_id	Any unique ID from the advertiser, such as loyalty membership ID, user ID, and external cookie ID.	a@example.com
Gender	ge	Single lowercase letter, f or m, if unknown, leave blank	f
Birthdate	db	Digits only with birth year, month, then day	19910526 for May 26, 1991.
City	ct	Lowercase with any spaces removed	menlopark
State or Province	st	Lowercase two-letter state or province code	ca
Zip or Postal Code	zp	String	94025
Country	country	Lowercase two-letter country code	us

Reference table from Meta's documentation describing the types of information that can be used for automatic Advanced Matching.

Event Data Details

The Meta Pixel event data parameters described either appear as URL parameters in GET requests or in the body of POST requests to the URL https://www.facebook.com/tr/. This appendix provides examples of the different

formats the Meta Pixel uses to send event data and annotates the parameters of interest for our study.

Acknowledgements

We thank Gunes Acar, assistant professor at the Digital Security group of Radboud University Nijmegen, and Steve Englehardt, privacy engineer at DuckDuckGo, for reviewing an earlier version of this methodology.

We would also like to thank the Mozilla Rally team, including Ted Han, Kim Bryant, Javaun Moradi, Jonathan Epstein, Robert Helmer, Bruce Etling, and Marnie Pasciuto-Wood.